

Please type a plus sign (+) inside this box →

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL

(Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))

Attorney Docket No. 730.37334X00

First Inventor or Application Identifier Rajeev KOODLI

Title See 1 in Addendum

Express Mail Label No. S-1103

PRO

S-1103

APPLICATION ELEMENTS

See MPEP chapter 600 concerning utility patent application contents.

1. * Fee Transmittal Form (e.g., PTO/SB/17)
(Submit an original and a duplicate for fee processing)
2. Specification [Total Pages 32]
(preferred arrangement set forth below)
 - Descriptive title of the Invention
 - Cross References to Related Applications
 - Statement Regarding Fed sponsored R & D
 - Reference to Microfiche Appendix
 - Background of the Invention
 - Brief Summary of the Invention
 - Brief Description of the Drawings (if filed)
 - Detailed Description
 - Claim(s)
 - Abstract of the Disclosure
3. Drawing(s) (35 U.S.C. 113) [Total Sheets 3]
4. Oath or Declaration [Total Pages]
 - a. Newly executed (original or copy)
 - b. Copy from a prior application (37 C.F.R. § 1.63(d))
(for continuation/divisional with Box 16 completed)
 - i. DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).

* NOTE FOR ITEMS 1 & 13 IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).

ADDRESS TO: Assistant Commissioner for Patents
Box Patent Application
Washington, DC 20231

5. Microfiche Computer Program (Appendix)
6. Nucleotide and/or Amino Acid Sequence Submission
(if applicable, all necessary)
 - a. Computer Readable Copy
 - b. Paper Copy (identical to computer copy)
 - c. Statement verifying identity of above copies

ACCOMPANYING APPLICATION PARTS

7. Assignment Papers (cover sheet & document(s))
8. 37 C.F.R. § 3.73(b) Statement Power of (when there is an assignee) Attorney
9. English Translation Document (if applicable)
10. Information Disclosure Statement (IDS)/PTO-1449 Copies of IDS Citations
11. Preliminary Amendment
12. Return Receipt Postcard (MPEP 503)
(Should be specifically itemized)
13. * Small Entity Statement(s) Statement filed in prior application (PTO/SB-09-12)
14. Certified Copy of Priority Document(s)
(if foreign priority is claimed)
15. Other:

16. If a CONTINUING APPLICATION, check appropriate box, and supply the requisite information below and in a preliminary amendment:

Continuation Divisional Continuation-in-part (CIP)

of prior application No. /

Prior application information: Examiner _____

Group / Art Unit: _____

For CONTINUATION or DIVISIONAL APPS only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS

<input checked="" type="checkbox"/> Customer Number or Bar Code Label	020457	or <input type="checkbox"/> Correspondence address below (Insert Customer No. or Attach bar code label here)
Name		
Address		
City	State	Zip Code
Country	Telephone	Fax

Name (Print/Type)

Carl I. Brundidge

Registration No. (Attorney/Agent)

29,621

Signature

Date

12/23/1999

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

Addendum

1. A SCHEME FOR DETERMINING TRANSPORT LEVEL INFORMATION IN THE PRESENCE OF IP SECURITY ENCRYPTION

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventors: Rajeev KOODLI and Senthil SENGODAN

Invention: A SCHEME FOR DETERMINING TRANSPORT LEVEL
INFORMATION IN THE PRESENCE OF IP SECURITY
ENCRYPTION

Antonelli, Terry, Stout & Kraus
Suite 1800
1300 North 17th Street
Arlington, VA 22209

SPECIFICATION

To All Whom It May Concern:

Be it known that we, Rajeev KOODLI and Senthil SENGODAN,
citizens of India, have invented new and useful improvements
in:

A SCHEME FOR DETERMINING TRANSPORT LEVEL INFORMATION IN THE
PRESENCE OF IP SECURITY ENCRYPTION

of which the following is a specification.

SCHEME FOR DETERMINING TRANSPORT LEVEL INFORMATION IN THE
PRESENCE OF IP SECURITY ENCRYPTION

5 BACKGROUND OF THE INVENTION

The present invention relates to a method and apparatus for permitting access to information contained in selected fields of a packet subjected to security processing, particularly by intermediate nodes between source and 10 destination nodes when the packet is transmitted on a packet switched network from the source node to the destination node.

In packet switched networks, packets are transmitted between nodes connected to the network to effect communication between the nodes. Information in the packets may include 15 messages and commands such as a request for service, connection management controls, or data. The packets transmitted on the packet switched network are transmitted according to, for example, the Internet Protocol (IP) which defines how to format various information in the packets for transmitting on the 20 packet switched network. Currently, the IP is defined according to IPv4 where the "v4" indicates version 4 of the Internet Protocol, and IPv6 where the "v6" indicates version 6 of the Internet Protocol.

Security across the packet switched network is 25 afforded by IP Security Protocol (IPsec) as set forth in "Security Architecture for the Internet Protocol", by S. Kent et al, Request for Comments (RFC) 2401, Networking Group, Internet Engineering Task Force (IETF), November 1998. IPsec

is a protocol designed to provide various security services for traffic at the IP layer and upper layer protocols in both the IPv4 and IPv6 environments.

These security services are implemented through the 5 use of two traffic security protocols, namely, Authentication Header (AH) Protocol, and Encapsulating Security Payload (ESP) Protocol and through the use of Cryptographic Key Management Procedures and Protocols. The set of IPsec Protocols employed in any context, and the ways in which they are employed, are 10 determined by the security and system requirements of users, applications and/or sites/organizations.

The AH protocol provides security services of connection list integrity, data origin authentication and anti-replay. The ESP protocol provides security services of 15 confidentiality (encryption), and limited traffic flow confidentiality. The ESP protocol can also provide security services of connection list integrity, data origin authentication and anti-replay. These protocols may be applied alone or in combination with each other to provide a desired 20 set of security services in IPv4 and IPv6. When either of these protocols are used, particularly security processings are performed. Each protocol supports two modes of use: Transport Mode and Tunnel Mode. In the Transport Mode the protocols provide protection primarily for upper layer protocols. In the 25 Tunnel Mode the protocols are applied to tunneled IP packets.

Both AH and ESP make use of Security Associations

(SAs). An SA is a simplex "connection" that affords security services to the traffic carried by it. An SA is uniquely identified by a triple including a destination IP address, a Security Protocol (AH or ESP), and a Security Parameter Index 5 (SPI). The endpoints of an SA can, for example, be a pair of hosts, a pair of security gateways or a security gateway and a host. The security gateway can, for example, be a router, a firewall, etc. When either end or both ends of an SA are a security gateways then the tunnel mode is used.

10 Prior to the use of the above described security protocols, an SA must be established. Thus, for example, prior to use of the ESP protocol between, for example, two hosts, an SA, identified according to the triple described above, must be established between the pair. When the ESP protocol is used, the entire IP packet, with the exception of the IP header, the AH protocol header if present and the ESP header, is encrypted 15 during packet transmission as illustrated by the shaded areas in Fig. 1.

As illustrated in Fig. 1, the IP packet which has 20 been subjected to security processings according to the ESP protocol includes an IP header 102, an ESP header 104, a transport PDU (Payload Data Unit) 106, an ESP trailer 108 and an ESP authenticator 110. The ESP header 104 includes an SPI field 104-1 and a sequence number field 104-2. The transport 25 PDU 106 includes payload data which can be of variable length, and the ESP trailer 108 includes a padding field 108-1, a

padding length field 108-2 and a next header identifier field 108-2. The payload data can include, for example, transport level information including Transmission Control Protocol (TCP) Header, User Datagram Protocol (UDP) Header or Internet Control 5 Message Protocol (ICMP) Header and port number. The sequence number contained in the sequence number field 104-2 of the ESP header is a 32-bit monotonically increasing number that is present to prevent replay attacks. Since the payload data 106 and the ESP trailer 108 are encrypted according to the security 10 processing effected by, for example, the ESP protocol the only information visible to any intermediate nodes between the source and destination of the SA (including the first-hop-ingress-egress routers) would be the IP header 102, the ESP header 104 and the ESP authenticator 110.

15 Although not illustrated, an IP packet subjected to security processings according to the AH protocol in addition to ESP, when used, causes the same problems as when the ESP protocol above is used. Namely, the entire transport payload including the transport headers is encrypted. Therefore, an IP 20 packet subjected to security processings according to the AH and ESP protocol would have a configuration similar to that illustrated in Fig. 1 with the exception that the AH header would be present as well.

25 The security processing, particularly encryption, of selected information in selected fields especially the transport level information including the transport protocol

and the port number, is undesirable for certain intermediate nodes such as, for example, packet classifier/marker nodes providing differentiated services, policing nodes such as firewalls, or management nodes for metering. For example, if 5 the ingress router also serves the function as a packet classifier/marker node providing differentiated services, this router will not be able to access the transport level information, such as transport protocol and port number that may be used as indicators for indicating the type of service 10 desired according to the differentiated services framework.

For example, a TCP destination port number of 80 may indicate access to a webserver, which could be used to classify and mark the Differentiated Services Code Point (DSCP) in the packet for an appropriate treatment for web related services. This type 15 of classification and marking known as multi-field classification (and marking) and differentiated services (Diffserv) would not be possible if ESP is used as described above.

Therefore, there is a need to provide selected field 20 accessibility for certain intermediate nodes in the presence of, for example, security processing of an IP packet without compromising the security of the system nor requiring modifications to the existing protocols.

SUMMARY OF THE INVENTION

The present invention provides a method and apparatus for permitting access to selected information in selected fields of a packet, which has been subjected to security processing, by intermediate nodes in a packet switched network when the packet is being transmitted on the packet switched network.

Particularly, the present invention provides a method and apparatus which permits access to selected information such as transport level information, normally included in a payload of a packet upon which encrypting security processing has been performed according to a security protocol by a node in a packet switched network during transmission of the packet.

The present invention accomplishes the above by prior to performing security processing according to an encrypting security protocol on a packet, storing information related to selected information normally included in a payload of the packet in a header of the packet. The field in the header that this information is placed is not subjected to the encrypting security processing. Thereafter, the security processing is performed on the packet. The packet including the header having stored therein information corresponding to the selected information normally included in the payload and the payload upon which security processing has been performed is then transmitted on the packet switched network to its destination.

Therefore, according to the present invention, by

storing information related to the selected information normally included in the payload of the packet in the header of the packet as well, access to the selected information by intermediate nodes between the source and destination nodes in 5 a packet switched network is possible. The information related to the selected information normally included in the payload of a packet can, for example, be representative of the selected information or the selected information itself. The information related to the selected information can be freely accessed by the intermediate nodes, decoded if necessary, and 10 used to initiate specified processings on the packet by the intermediate node.

The present invention is particularly applicable to the security services provided by the IP Security (Ipsec) 15 Protocol, wherein the Encapsulating Security Payload (ESP) protocol is used. Implementing ESP on an IP packet requires that portions of the IP packet, particularly the payload of the IP packet, which includes transport level information such as transport protocol and port number, be encrypted during 20 transmission. Encrypting portions of the IP packet, particularly the payload of the IP packet eliminates access by intermediate nodes to the transport level information contained in the encrypted portions of the IP packet when the IP packet is transmitted between source and destination nodes on a packet 25 switched network.

The present invention allows for access of the

transport level information contained in the encrypted portions of the IP packet during transmission of the IP packet by storing information related to the transport level information in the security protocol header of the IP packet. The 5 information selected to the transport level information could, for example, be stored in the sequence number field of the security protocol header of the IP packet. Intermediate nodes in the packet switched network can easily access such information since it is part of the security protocol header 10 which is not encrypted according to security processings effected on the IP packet according to the encrypting security protocol.

BRIEF DESCRIPTION OF THE DRAWINGS

15 The present invention will be more apparent from the following detailed description, when taken in conjunction with the accompanying drawings, in which:

Fig. 1 is a schematic diagram illustrating the configuration of an IP packet when the ESP protocol is used;

20 Figs. 2a-c are schematic diagrams illustrating different scenarios for Security Associations;

Fig. 3 is a schematic diagram illustrating the sequence number field of the IP packet when the ESP protocol is used according the present invention;

25 Figs. 4a and b are tables illustrating the values of various bits of the sequence number field of an IP packet

configured according to the present invention; and

Figs. 5a-c are schematic diagrams illustrating various modes of operation relative to the ESP protocol according the present invention.

5

DETAILED DESCRIPTION OF THE INVENTION

The present invention will be described relative to Figs. 1, 2a-c, 3, 4a-b, and 5a-c. It should noted that the present invention is not limited to the embodiments illustrated in these figures and other embodiments recognizable by those of ordinary skill in the art are encompassed within the scope of the present invention.

The present invention is intended to operate in a packet switched network where security services according to IPsec ESP are provided. These security services effected by security processings performed according to the ESP protocol of the IPsec is initiated once a Security Association (SA) has been established. A SA is a simplex "connection" that affords security services to the traffic carried by it. A SA is uniquely identified by a triple including a destination IP address, security protocol (AH or ESP), and a Security Parameter Index (SPI). The endpoints of a SA can, for example, be a pair of hosts, a pair of security gateways or a security gateway and a host. The security gateway can, for example, be a router, a firewall, etc.

Figs. 2a-c illustrate examples of scenarios for the

PCT/US2003/033330

establishment of SAs. Each figure illustrates an SA including two sources, source 1 (S1), 201 and source 2 (S2), 203, one destination 205 and 3 SAs 207. Each SA 207 of each figure terminates at the same destination 205.

5 As illustrated in Fig. 2a, the SAs established between S1 201 and the destination 205 and S2 203 and the destination 205 are all identical as per the solid line arrows. Particularly, between S1 201 and destination 205 two identical SAs 207-A are established and between S2 203 and destination 10 205 a single SA 207-A identical, as per the solid line arrow, to the SAs established between S1 201 and destination 205 is established.

15 Fig. 2b illustrates a case where the two SAs 207-A established between S1 201 and destination 205 are identical as per the solid line arrows. An SA 207-B established between S2 203 and destination 205 is different as per the dashed line arrow, from the SAs 207-A established between the S1 201 and the destination 205.

Fig. 2c illustrates a case where all SAs established 20 between the sources S1 201 and S2 203 and the destination 205 are different from each other. Namely, between S1 201 and the destination 205 an SA 207-A is established as per the solid line arrow and a different SA 207-C is established as per the dotted line arrow. Further, between the S2 203 and the 25 destination 205, a SA 207-B different from the SA 207-A and SA 207-C, as per the dashed line arrow is established.

As per Figs. 2a-c, the choice of identifying an SA lies with the receiver (destination). Since the SA is uniquely identified by the triple destination IP address, security protocol and SPI, the destination is required to consider certain issues while determining whether a SA is unique or distinct. For example, when the security protocol (AH, ESP) is different, then the SA established between the source and destination is different. However, two SAs that terminate at the same destination and use different security protocols can, for example, use identical SPIs. Further, two SAs that terminate at the same destination and which use the same security protocol may be either identical or distinct. In this situation, the SA would be identical if the respective SPI values are identical. Otherwise, the SAs in this situation would be different. A destination must use different SAs if the security processing namely, encryption algorithm, key links, key value, etc., is different. If the security processing is identical then the destination may use identical SAs. However, the destination may use different values for the SA if it so desires.

In the conventional technique, once an SA has been established and security processings according to the ESP protocol (and optionally AH) has been initiated, the IP packet being transmitted on the packet switched network is entirely encrypted with the exception of portions or all of the IP header and the security protocol (AH or ESP) header. Such a

packet is illustrated, for example, in Fig. 1 relative to the
ESP protocol. As shown in Fig. 1, the transport PDU 106 and
the ESP trailer 108 are fully encrypted whereas the IP header
102, the ESP header 104 and the ESP authenticator 110 are not
5 encrypted. The encrypted areas are shown by the shading.

The conventional technique as described above,
suffers from numerous disadvantages one of which being that
selected information in selected fields of the IP packet are
encrypted, thereby causing such information to be inaccessible
10 to intermediate nodes which may use such selected information
to transmit or perform other processings on the packet. This
now inaccessible information could, for example, be transport
level information which is used by the intermediate nodes to
provide, for example, differentiated services, policing or
metering. The transport level information includes, for
15 example, transport protocol information including TCP, UDP and
ICMP and port number information. The transport level
information is normally carried in the payload data portion of
an IP packet. The inaccessibility of such information normally
20 carried in the payload portion of an IP packet does not permit
the intermediate nodes between the source and destination nodes
in a packet switched network to perform such value added
services relative to the communication.

The present invention overcomes the above by
25 permitting access to selected information normally included in
the payload of a packet upon which encrypting security

processing has being performed. Such is accomplished in the present invention by, for example, placing information related to the selected information in the security protocol header of the IP packet. This operation is conducted prior to security processing being performed on the packet according to the security protocol being used. The information related to the selected information could, for example, be representative of the selected information or the selected information itself.

5 The selected information could, for example, be transport level information including transport protocol and port number information. The information related to the selected information stored in the security protocol header is stored along with information identifying the security protocol.

10 After the above described storing operation, security processing is performed on the packet and the packet is transmitted to its destination on the packet switched network.

15 Since information related to the selected information is now contained in the security protocol header of the packet which is not encrypted, such information is readily accessible by intermediate nodes in the packet switched network existing

20 between source and destination nodes.

According to the present invention, the sequence number field of the security protocol header can be used to convey information related to the selected information when security processing has been performed and the packet is being transmitted. Other fields of the security protocol header or

the IP header could possible be used. The present invention is not limited to use of the sequence number field. The present invention provides a modified security protocol header specifically modifying the sequence number field 104-2. The 5 sequence number field 104-2 is modified to include information related to the selected information normally contained in the encrypted payload portion of the packet when security processing has been performed. The sequence number field 104-2 of the security protocol header, particularly the ESP header 10 illustrated in Fig. 1, is 32 bits long and can be configured such as that illustrated in Fig. 3.

As illustrated in Fig. 3, the leading 8 bits of the sequence number field 104-2 could be used to provide information related to the transport level information of the packet with the remaining 24 bits being used, for example, to provide information regarding the sequence number. While the specific number of bits of the sequence number field assigned to be used to provide information related to the transport 15 level information may vary, it is important that knowledge of the scheme used be provided to each of the intermediate nodes 20 of an established SA.

Specifically, an intermediate node that requires access to such information related to the transport level 25 information needs to be aware of the mechanism employed by the source node of the SA to represent such information in the sequence number field of the security processing header of the

packet. Several mechanisms are possible by which an intermediate node may be informed of the specific number and uses of the leading bits of the sequence number field to represent the transport level information. These mechanisms 5 can include, for example, standardizing the number and usage of the bits so that any intermediate node need merely to refer to predetermined (protocol) standard interpretations. Also the source node of the SA may indicate the number and usage of these bits to an intermediate node by communications which can 10 be performed on or off line. The present invention as illustrated, for example, in Fig. 3 provides a particular mechanism. However, the present invention is not limited to such a mechanism.

In Fig. 3, a Transport Level ID (TLID) 401 is provided. The TLID 401 includes a plurality of leading bits of the sequence number field 104-2. The TLID 401 includes a Transport Protocol ID (TPID) 403 and a Port Number ID (PNID) 405. The TPID 403 and the PNID 405 are each formed by a plurality of bits, the total of which are equal to the bits 15 used for the TLID 401. In the present invention as illustrated in Fig. 4, the TPID 403 comprises two bits (a, b) and the PNID 405 comprises six bits (c-h).

The TPID 403 identifies the particular transport protocol (TCP, UDP, ICMP, etc.) used and the PNID identifies the 20 associated port number. It should be noted that the exact 25 number of bits used for the each of the TPID 403 and PNID 405

can be different from that illustrated in Fig. 3. The present invention is not limited to the embodiment described above.

The present invention provides a Transport ID Mapping Table 500 as illustrated in Fig. 4a providing information indicating corresponding relations between the TPID 403 and the transport protocol 501. In the table, the TPID 403, including bits a and b, identifies the particular transport protocol 501 which forms part of the transport level information. The transport protocol 501 could, for example, include ICMP 501-1 corresponding to TPID bits "00" 403-1, UDP 501-2 corresponding to TPID bits "01" 403-2, TCP 501-3 corresponding to TPID bits "10" 403-3 and undefined 501-4 corresponding to TPID bits "11" 403-4.

The present invention also provides a Port Number ID mapping table 502 as illustrated in Fig. 4b providing information indicating corresponding relations between the PNID 405 and port numbers. In the table, the PNID 405 can be used to identify either a source port number 503, a destination port number 504 or both. A flag A 506 is provided to indicate whether a particular bit sequence of the PNID 405 has been assigned to a flow of packets (A=1) or not to any flow of packets (A=0).

The tables illustrated in Figs. 5 a and b can be provided, for example, by the source node of the SA to the intermediate nodes existing between the source node of the SA and the destination node of the SA. Any of the mechanisms

described above for accomplishing this can also be used. These tables are stored in the intermediate nodes for use in decoding the information conveyed by the leading bits of the sequence number field of the security protocol header of an IP packet.

5 The mapping of the values illustrated as being contained in these tables, as provided to any one of the intermediate nodes, is performed, for example, by the source node of a SA.

The mapping of the values of the tables illustrated in Figs. 4a and b can, for example, be unique for a triple including a source IP address, a destination IP address and an SPI. Inclusion of the source IP address in this triple permits the creation of a mapping locally within the source of an SA without the need for consulting other sources that may share the SA. It should be noted that SPI in the above noted triple corresponds to the SPI in, for example, the ESP header but not to the SPI in, for example, the AH header. The SPI of the AH header may also be present in the same IP packet. In other words, the security protocol field is implicit in the triple with a constant value of ESP. When a new six bit value needs 20 to be assigned to a new flow of packets, the source of the SA may utilize any available value between 0 and 63.

Since the mapping is unique to a triple including a source IP address, a destination IP address and an SPI, the session/flow is uniquely identified by a triple including the 25 source IP address, the destination IP address, the SPI and the PNID. Here the PNID refers to a portion of the leading bits of

the sequence number field of a packet configured according to the present invention.

It is possible that if the same SA is used for several simultaneous sessions between the same source and destination pair, the Port Number ID mapping table 502 could be exhausted, for example, after 64 UDP or 64 TCP sessions. In such a case, there would not be any free PNIDs that the source node could allocate to a new session. In order to alleviate this problem, the destination node of the SA could assist by changing the SA when the number of simultaneous TCP (or UDP) sessions from the same source node that require identical ESP processing equals 64. Note that this is an implementation issue and does not violate the IPsec protocol. It should be noted that the larger the number of bits that are used for this purpose, the faster the SA refresh is required to be performed for the particular SA.

Once the source node of the SA has performed the mapping, it must notify intermediate nodes that need transport level information of such mapped values. Specifically, information organized in a manner similar to the tables illustrated in Figs. 4a and b are provided to the intermediate nodes. The network points that need to be aware of such a mapped values include, for example: (1) a diff-serv marker node (2) a management metering node, or (3) a policy enforcement node (firewall). For a diff-serv marker node, the only node that needs to be notified is the node that is responsible for

marking the packets that originate from the source node. For a policy enforcement node or firewall, it is likely that all firewalls that need to police the traffic need to be informed of the mapped values. Similarly, for a management metering node, all nodes that are interested in metering the traffic need to be aware of the mapped values. The specific mechanisms that the source node of the SA employs to deliver the mapped values to intermediate nodes can be of any type as described above including those that may become known to one skilled in the art.

Figs. 5a-c are schematic diagrams illustrating various modes of operation relative to the ESP protocol according to the present invention. The diagrams illustrated in Figs. 5a-c are provided as examples of how the present invention operates in a packet switched system. Although the ESP protocol alone is used in the schematic diagrams illustrated in Figs. 5a-C, other security protocols used along with ESP, particularly the AH protocol used with ESP would operate in a similar manner.

Fig. 5a illustrates a configuration where the ESP protocol is operated in the transport mode, and where source 1 (S1) 601 and source 2 (S2) 602 and destination (d) are sources and destinations of various SAs. Further, gateway (G1) 607 is an egress router and gateway (G2) 609 is an ingress router. First, a case will be considered where two UDP sessions are in existence between S1 601 and D 605, and one TCP session is in

existence between S2 602 and D 605. In this case, it is assumed that the ESP protocol with the same algorithms, keys and other security parameters is desired for each of the three sessions. Since the number of simultaneous sessions using the 5 same transport protocol (TCP/UDP) and requiring identical ESP protocol security processing (encryption algorithm, keys etc.) that originates at any single host and terminates at D 605 does not exceed 64, D may use the same SA for each of these sessions. The common SA for each of these three sessions is 10 identified by the triple <D's IP address, ESP, SPI>, where SPI is a specific value for the SPI. In this establishment of SAs using the Internet Key Exchange (IKE), there does not exist a mechanism whereby the transport protocol (TCP/UDP) information is conveyed to the destination of the SA at the time of SA 15 establishment. This then implies that one of two things may be done:

20 (1) Using the existing IKE mechanisms, the destination D 605 may use a common SA for a maximum of 64 simultaneous sessions originating from the same source (and terminating at D 605) which require identical ESP protocol security processing. This is independent of whether these sessions are UDP or TCP.

25 (2) The existing IKE mechanisms may be extended so that the destination D 605 is aware of the transport protocol of the session at the time of SA establishment. In such a case, the destination D 605 may use a common SA for a maximum of 64 simultaneous sessions originating from the same source (and

terminating at D 605) which require identical ESP protocol security processing and which also use the same transport protocol (TCP/UDP).

Now assume that, for each of the three sessions,

5 transport level information needs to be sent to the egress router G1 607 for marking purposes and the ingress router G2 607 for policing/metering purposes. Source S1 601 assigns an available TLID to each of the UDP sessions, where the leading two bits of the TLID are set to "01" (indicating UDP).

10 Similarly, source S2 602 assigns an available TLID to the TCP session terminating at D 605 such that the leading two bits of the TLID are set to "10" (indicating TCP). These sessions are identified by the triple <(source) IP address, D's IP address, SPI, TLID>. Since the destination D 605 has ensured that not more than 64 simultaneous sessions (TCP & UDP combined with existing IKE; TCP or UDP with modified IKE) with the same ESP protocol security processing share the same SA, it is guaranteed that a TLID will be available at the source for assignment.

20 Fig. 5b illustrates a configuration where the ESP protocol is operated in the usage tunnel mode. In Fig. 5b, S1 601 and S2 602 are session sources and D 605 is the common destination. The sessions (from S1 601 and S2 602) pass through firewalls F1 611 and F2 613, as well as through G1 607 and G2 609. An SA exists between F1 611 and F2 613 while G1 607 and G2 609 (intermediate nodes) need transport level

information for each of the sessions between S1 601 or S2 602 and D 605 for the purpose of packet marking or management metering.

First, a case is considered where two UDP sessions exist between S1 601 and D 605, and a TCP session exists between S2 602 and D 605. Packets from each of the three sessions flow through F1 611, G1 607, G2 609 and F2 613 before going on to D 605. Packets belonging to each of the three session are afforded the same ESP protocol security processing.

In other words, the same ESP protocol security processing for the SA between established F1 611 and F2 613 is used for packets belonging to each of the three sessions.

The SA between F1 611 and F2 613 is uniquely identified by the triple <F2's IP address, ESP, SPI>. The conditions for the destination F2 613 of the SA to change the SA for identical ESP protocol security processings are as follows:

(1) Using the existing IKE mechanisms, the destination F2 613 of the SA may use a common SA for a maximum of 64 simultaneous sessions originating from F1 611 (and terminating at F2 613) which require identical ESP protocol security processing. This is independent of whether these sessions are UDP or TCP sessions and is also independent of the source (such as S1 601 and S2 602) that these sessions originate from.

(2) The existing IKE mechanisms may be extended so

that the destination F2 613 of the SA is aware of the transport protocol of the session at the time of SA establishment. In such a case, the destination F2 613 may use a common SA for a maximum of 64 simultaneous sessions originating from the same 5 F1 611 (and terminating at F2 613) which require identical ESP protocol security processings and which also use the same transport protocol (TCP/UDP).

(3) The existing IKE mechanisms may be further extended so that the destination F2 613 of the SA is aware of 10 the session source (S1 601 or S2 602) at the time of SA establishment. In such a case, the destination F2 613 may use a common SA for a maximum of 64 simultaneous sessions originating from the same F1 611 (and terminating at F2 613) which require identical ESP protocol security processing, which 15 use the same transport protocol (TCP/UDP) and which originate at the same source. This is used only when the session source is also needed by the intermediate node. Since the session source IP address is encrypted in the tunnel mode of the ESP protocol, and an intermediate node desires this information, 20 then such information needs to be suitable mapped within the sequence number field or by other means.

After information related to the transport level information is stored in the sequence number field of the ESP protocol header, F1 611, then conveys the mapping tables to the 25 intermediate nodes G1 607 and G2 609 that requires such transport level information to permit decoding of such

information.

Fig. 5c illustrates a case involving a nested SA scenario. As seen in Fig. 5c, an end-to-end transport-mode ESP protocol SA exists between the session source (S1 601 or S2 602) and the session destination (D 605). In addition, this SA is nested within a tunnel mode ESP protocol SA existing between F1 611 and F2 613. In this case, destination D 605 determines whether a new SA should be created for sessions requiring identical ESP protocol security processing based on the conditions discussed for the case illustrated in Fig. 5a. In this case, however, the session sources pass on the mapping tables to the source of a first security gateway existing in F1 611, and not to the intermediate nodes (G1 607 and G2 609) that require transport level information. In addition, the mapping tables are passed on to any other intermediate nodes (such as G3 615) that are present beyond the termination of the nested SA. F2 613 determines whether a new SA needs to be created between F1 611 and F2 613 based on the conditions discussed in the case illustrated in Fig. 5b. F2 613 creates new mapping tables which are used to configure the bits of the leading byte of the sequence number field to represent the transport level information (protocol and port number). This is done in a similar fashion to that described in the case illustrated in Fig. 5b.

When a packet arrives at F1 611 from S1 601, F1 611 examines the leading byte of the sequence number field and

based on the mapping tables that it received from S1 601, F1 611 determines the transport level information. Using mapping tables created by F1 611, F1 611 then inserts suitable values into the leading byte of the sequence number field of the outer 5 ESP protocol header. At F2 613, the outer ESP protocol header is removed, and the transport mode ESP protocol packet is what flows between F2 613 and D 605.

The present invention as described above provides a mechanism by which transport level information can be conveyed 10 to intermediate nodes in the presence of IPsec security protocols, particularly ESP protocol and AH protocol. The transport level information of importance here are the transport protocol (TCP/UDP/ICMP) and the port number. The mechanism was discussed for transport mode, tunnel mode and nested ESP protocol security processings. Using such mechanisms, intermediate nodes such as diff-serv 15 markers/classifiers, firewalls/policers as well as network management nodes can perform their respective functions. These nodes require such transport level information for their 20 functioning.

While the present invention has been described in detail and pictorially in the accompanying drawings it is not limited to such details since many changes and modifications recognizable to those of ordinary skill in the art may be made 25 to the invention without departing from the spirit and the scope thereof.

WE CLAIM:

1. A method of permitting access to selected information normally included in a payload of a packet upon which encrypting security processing has been performed by a node in a packet switched network during transmission of the packet, said method comprising:

prior to performing encrypting security processing on a payload of a packet, storing information corresponding to selected information normally included in a payload of said packet in a field in a header of said packet where said field is not subject to said encrypting security processing;

performing said encrypting security processing on said payload of said packet; and

transmitting said packet including said header and said payload upon which encrypting security processing has been performed in the packet switched network, thereby permitting access to said selected information normally included in said payload of said packet via said header of said packet by a node in the packet switched network.

2. A method according to claim 1, wherein said selected information includes transport level information.

3. A method according claim 2, wherein said transport level information includes transport protocol information

including Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP) and port number information.

4. A method according to claim 1, wherein said selected information is stored in a security protocol header of said header of said packet, said security protocol header not being subject to said encrypting security processing.

5. A method according to claim 4, wherein said selected information is stored in a sequence number field of said security protocol header of said packet.

6. A method according to claim 5, wherein said sequence number field is modified to include information representative of said selected information.

7. A method according to claim 6, wherein leading bits of said sequence number field are used to provide information related to transport level information of said packet and remaining bits of said sequence number field are used to provide information regarding a sequence number of said packet.

8. A method according to claim 7, wherein said transport level information includes a transport protocol ID identifying a particular transport protocol to be used and a port number ID

identifying a particular port to be used.

9. A method according to claim 1, wherein said encrypting security processing is performed according to the Encapsulating Security Payload (ESP) protocol.

10. A method according to claim 1, wherein said encrypting security processing is performed according to the Authentication Header (AH) protocol.

11. A packet switched network comprising:

 a network; and
 a plurality of nodes interconnected to each other by said network to permit communication between said nodes using packets,

 wherein each node transmits a packet to another node in a manner to permit access to selected information normally included in a payload of said packet upon which encrypting security processing has been performed during said transmission of said packet from said node to said another node,

 wherein node, prior to performing encrypting security processing on a payload of a packet, stores information corresponding to selected information normally included in a payload of said packet in a field in a header of said packet where said field is not subject to said encrypting security processing, performs said encrypting security processing on

said payload of said packet, and transmits said packet including said header and said payload upon which encrypting security processing has been performed, thereby permitting access to said selected information normally included in said payload of said packet via said header of said packet during transmission of said packet from said node to said another node.

12. A packet switched network according to claim 11, wherein said selected information includes transport level information.

13. A packet switched network according claim 12, wherein said transport level information includes transport protocol information including Transmission Control Protocol (TCP), User Datagram Protocol (UDP), Internet Control Message Protocol (ICMP) and port number information.

14. A packet switched network according to claim 11, wherein said selected information is stored in a security protocol header of said header of said packet, said security protocol header not being subject to said encrypting security processing.

15. A packet switched network according to claim 14, wherein said selected information is stored in a sequence

number field of said security protocol header of said packet.

16. A packet switched network according to claim 15, wherein said sequence number field is modified to include information representative of said selected information.

17. A packet switched network according to claim 16, wherein leading bits of said sequence number field are used to provide information related to transport level information of said packet and remaining bits of said sequence number field are used to provide information regarding a sequence number of said packet.

18. A packet switched network according to claim 17, wherein said transport level information includes a transport protocol ID identifying a particular transport protocol to be used and a port number ID identifying a particular port to be used.

19. A packet switched network according to claim 11, wherein said encrypting security processing is performed according to the Encapsulating Security Payload (ESP) protocol.

20. A packet switched network according to claim 11, wherein said encrypting security processing is performed according to the Authentication Header (AH) protocol.

ABSTRACT

A method and apparatus which permits access, by intermediate nodes between source and destination nodes, to selected information such as transport level information, normally included in a payload of a packet upon which encrypting security processing has been performed according to an encrypting security protocol. In the present invention, prior to performing encrypting security processing on the packet, according to the security protocol, information related to selected information normally included in a payload of the packet is stored in a field in the header of the packet where the field is not subject to the encrypting security processing. Thereafter, encrypting security processing according to the security protocol is performed on the packet. The packet including the header having stored therein information corresponding to the selected information normally included in the payload and the payload upon which encrypting security processing has been performed is then transmitted on the packet switched network to its destination. Since the information related to the selected information normally included in the payload of the packet is stored in the header of the packet, access to the selected information by the intermediate nodes between source and destination nodes in a packet switched network is possible.




FIG 1




FIG 2A

FIG 2B

FIG 2C




FIG 3

4 ⁰³		5 ⁰¹	5 ⁰⁰
a	b	Transport Protocol	
403-1	00	ICMP	501-1
403-2	01	UDP	501-2
403-3	10	TCP	501-3
403-4	11	Undefined	501-4

FIG 4 A

c ⁰⁵ d ⁰⁶ e ⁰⁷ f ⁰⁸ g ⁰⁹ h ⁰⁰	A	S_port	D_port	5 ⁰⁴ 5 ⁰²
000000	1	1050	857	
000001	1	880	857	
000010	1	2240	80	
000011	0			
000100	1	1769	21	
000101	0			
000110	0			
111110	1	78	57	
111111	1	136	753	

FIG 4 B

FIG 5A




FIG 5B




FIG 5C

